

Deciding what to display: maximizing the information value of social media

Sandra^{*†}
Servia-Rodríguez
AtlantTIC, University of Vigo
sandra@det.uvigo.es

Bernardo A. Huberman
HP Labs
bernardo.huberman@hp.com

Sitaram Asur
HP Labs
sitaram.asur@hp.com

ABSTRACT

In information-rich environments, the competition for users' attention leads to a flood of content from which people often find hard to sort out the most relevant and useful pieces. Using Twitter as a case study, we applied an attention economy solution to generate the most informative tweets for its users. By considering the novelty and popularity of tweets as objective measures of their relevance and utility, we used the *Huberman-Wu* algorithm to automatically select the ones that will receive the most attention in the next time interval. Their predicted popularity was confirmed by using Twitter data collected for a period of 2 months.

1. INTRODUCTION

The popularity of the Web and social media services has resulted in a constant flood of information which makes it hard for users to identify and consume the most relevant and useful pieces of content. Given the limited amount of attention that users can afford, providers of content have to decide what items to prioritize in order to gain the attention of users and become popular.

In earlier research, the task of automatically selecting the most relevant and useful pieces of information has been approached from different perspectives. *Ranking* [2] is at the core of the Information Retrieval (IR) scientific discipline behind search engines as *Google* or *Yahoo! Search*, where pieces of information -documents- are ranked according to their relevance to a given query. *Recommendation* [18], the discipline behind the success of many online services as shopping, photo(music, video)-sharing or online social networks, aims to predict the pieces of information that users will find more useful; either because these pieces are (i) similar to previous pieces liked by them -*content-based filtering*-, or (ii) liked by other users with similar preferences -*collaborative*

filtering. However, neither ranking nor recommendation are suitable for deciding the content to prioritize in social media, since the former requires a query to answer and the latter the preferences of the subjects (readers) to receive the recommendation. So far, in online newspapers, magazines and blogs, editors have been the ones to decide the choice of content and the presentation order. However, the emergence of news media aggregators, such as *digg.com* or *reddit.com*, has led to a citizen journalism-based ordering. That is, instead of having professional editors to determine the important news, people vote for news that they find interesting and the votes received by an article play an important role in its ranking with respect to other news on the front page or in the different ordered lists of news.

Social media services such as Twitter feature a large number of subscribers and they serve as aggregators of content such as news, promotional campaigns, media and status updates from users. Given the diversity and magnitude of content that is available, it is important, from the service provider's point of view to ensure easy access to relevant information to users, in order to retain and increase user engagement with the platform.

The Twitter timeline displays all tweets from the users that subscribers follow in decreasing order of publication. However, novelty is not the only feature that makes tweets valuable to users. Other features such as popularity, can also contribute to give value to tweets. Although Twitter has recently included a feature to display in bigger size those tweets that have received more engagement (retweets, favourites,...), users still have to scroll down their timeline in order to find interesting tweets. We are convinced that, apart from novelty, popularity should play a key role in deciding which tweets should be displayed in the top positions in order to increase users' engagement with the platform. This would specially be useful for those users who log onto Twitter using mobile devices (mobile phones, tablets,...), due to the reduced dimension of their screens (when compared with computers).

Within this setting, we study a method for selecting the optimal arrangement of tweets that maximizes the information value of users. Considering the number of retweets as an indicator of the popularity of a tweet and the time since it was posted as an indicator of its novelty, we conduct a study to empirically validate the solution proposed by Huberman and Wu in [12] to obtain the optimal arrangement of tweets that maximizes the informativeness for the users. By mapping the problem to that of optimal allocation of effort for a number of competing projects, Huberman and Wu for-

^{*}Corresponding author.

[†]This work was partially done while Sandra Servia-Rodríguez was an intern at HP Labs.

mulate the problem as a special case of *bandit problem* [10, 20], which they finally solve by applying the adaptive greedy algorithm proposed by Bertsimas and Niño-Mora [3].

To evaluate our arrangement method, we crawled the Twitter streams of 5 different influential news media accounts (*New York Times*, *BBC*, *CNN*, *Huffington Post* and *Mashable*) and collected their tweets and retweets for a period of two months. We obtained an arrangement taking into account the popularity and novelty of content and validated this using the actual attention provided by users in the form of retweets, replies and favorites. Our results show high accuracy of prediction of user attention thus demonstrating the benefits of our proposed solution. We believe that these findings will be very useful for content providers to automatically organize the most informative items for their customers and, in this way, win their attention.

2. BACKGROUND

As mentioned, our goal with this research is to demonstrate the suitability of using the approach in [12] to the problem of selecting the most informative tweets to users. This solution involves the steps of (i) mapping the problem of optimizing the information one gets to that of the optimal allocation of effort to a number of competing problems, (ii) formulating the problems as a special type of *bandit problem*, *dual-speed restless bandit problem*, and (iii), by using the adaptive greedy algorithm developed by Bertsimas and Niño-Mora [3], calculate an index for each item's state, which is then used to decide which item goes into the top list of a given time. Below we detail these steps.

2.1 Problem formulation

Consider a system that wishes to present n items to a large group of users but it can only present k ($k < n$) at any given time. Since an item displayed in front of a user has a higher probability of being chosen than when it is not displayed, we will call these k items the “top list”. We will also assume that the system can update its top list at discrete times $t = 0, 1, 2, \dots$

If the system can track a certain set of properties for each item, such as its reputation, history or age, we say that the item is in a “state” defined by those properties. Let E be the set of all possible states, i.e. all possible combination of those trackable properties. In general, the state of an item may change as time goes on. We assume that the state of each item changes according to a Markov process independent of the state of other items, with transition probabilities $\{P_{ij}^1 : i, j \in E\}$ if the item is on the top list, and $\{P_{ij}^0 : i, j \in E\}$ if it is not. We also make the assumption that having an item on the top list encourages more users to try it out and consequently accelerates its transitions from one state to the other. Conversely, when an item transitions away from the top list it slows down its rate of change by an amount ϵ_i which is less than 1. This *dual-speed assumption* can be stated as

$$P_{ij}^0 = \begin{cases} \epsilon_i P_{ij}^1, & i \neq j \\ (1 - \epsilon_i) + \epsilon_i P_{ii}^1, & i = j \end{cases} \quad (1)$$

where $\epsilon_i \in [0, 1]$.

Consider the total expected utility r_i obtained at one time step by those users who decide to try an item on the top list which has state i . This utility may depend on many

factors, such as the total expected number of users choosing the item at a given time step, or the expected quality of the item. Since we can always enlarge the definition of “state” to include these factors, the utility r_i is uniquely determined by the item state i . In other words, we can assume that $r = (r_i)_{i \in E}$ is an $|E|$ -dimensional constant vector known by the system.

Our goal then is to design a system that maximizes the total expected utility of all users:

$$\max_{u \in U} \mathbb{E}_u \left[\sum_{t=0}^{\infty} \sum_{m=1}^n \beta^t r_{i_m(t)} I_m(t) \right], \quad (2)$$

where $0 < \beta \leq 1$ is the future discount factor, $i_m(t)$ is the state of item m at time t , and

$$I_m(t) = \begin{cases} 1 & \text{if item } m \text{ is displayed at time } t, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

We seek to find the optimal strategy, u , in the space U of stationary strategies (strategies that depend on current item states only). This optimal strategy can then get translated into the set of offerings that should appear in the top list.

2.2 Solution

The model described is a *dual-speed restless bandit* problem: *restless* because changes of state can also occur when the items are not displayed in the top list and *dual-speed* because those changes do happen at a different speed than those on the top list. Bertsimas and Niño-Mora showed that a *relaxed* version of the dual-speed problem is always indexable—it is possible to attach an *index* to each item state, so that the top list is the one including those items with the largest indices—and proposed an efficient adaptive greedy algorithm to compute these indices [3].

Before using the algorithm, it is necessary to calculate a set of constants A_i^S . Assuming that E is finite, for any subset $S \subseteq E$, we define the S -active policy u_S to be the strategy that recommends all items whose state is in S . Considering an item that starts from an initial state $X(0) = i$, under the action implied by strategy u_S , its total occupancy time in S is given by

$$V_i^S = \mathbb{E}_{u_S} \left[\sum_{t=0}^{\infty} \beta^t I_S(t) | X(0) = i \right], \quad (4)$$

where

$$I_S(t) = \begin{cases} 1 & \text{if } X(t) \in S, \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

We have

$$V_i^S = \begin{cases} 1 + \beta \sum_{j \in E} P_{ij}^1 V_j^S, & i \in S, \\ \beta \sum_{j \in E} P_{ij}^0 V_j^S, & i \in S^c. \end{cases} \quad (6)$$

The variables $\{V_i^S\}_{i \in E}$ can be solved from the set of linear equations above.

A matrix of constants $\{A_i^S\}_{i \in E, S \subseteq E}$ is defined by means of V_i^S , as follows:

$$A_i^S = 1 + \beta \sum_{j \in E} P_{ij}^1 V_j^{Sc} - \beta \sum_{j \in E} P_{ij}^0 V_j^{Sc}. \quad (7)$$

The constants $\{A_i^S\}$ are then used in the Bertsimas-Niño-Mora algorithm as indicated in Algorithm 1.

Algorithm 1 Bertsimas-Niño-Mora adaptive greedy algorithm

Step 1. Set $S_{|E|} = E$ and

$$y^{S_{|E|}} = \max \left\{ \frac{r_i}{A_i^E} : i \in E \right\}. \quad (8)$$

Select $\pi_{|E|}$ as any maximizer and set $G_{\pi_{|E|}} = y^{S_{|E|}}$.

Step 2. For $k = 2, 3, \dots, |E|$, set $S_{|E|-k+2} \setminus \{\pi_{|E|-k+2}\}$ and

$$y^{S_{|E|-k+1}} = \max \left\{ \frac{r_i - \sum_{j=1}^{k-1} A_i^{S_{|E|-j+1}} y^{S_{|E|-j+1}}}{A_i^{S_{|E|-k+1}}} : i \in S_{|E|-k+1} \right\}. \quad (9)$$

Select $\pi_{|E|-k+1}$ as any maximizer and set $G_{\pi_{|E|-k+1}} = G_{\pi_{|E|-k+2}} + y^{S_{|E|-k+1}}$.

Finally, the strategy is to always display the k items whose states have the largest G index.

3. DECIDING WHAT TO DISPLAY ON TWITTER

On Twitter, the home timeline is a long stream showing all tweets from those that users have chosen to follow, displayed, at a given time, in decreasing order of publication. Focusing on a particular group of users, influential news media, we are interested on *selecting the optimal arrangement of tweets to be displayed to their followers* in order to maximize their informational value and, in this way, grab users' attention.

With this aim, we particularize the *Huberman-Wu* algorithm (Section 2) to this scenario. In order to properly define the states and compute the transition probabilities between these states, we use actual data from Twitter. Specifically, we have been monitoring the Twitter accounts of 5 different news media and the retweets done to their tweets for a period of 2 months. In what follows, we describe the dataset obtained and how we use the tweets of the first month to set up the states and transition probabilities and the data of the last month to prove the suitability of our approach (Section 4).

3.1 Dataset

Our dataset consists of tweets and retweets collected from crawling the Twitter streams of 5 different influential news media accounts (*The New York Times*, *BBC Breaking News*, *CNN Breaking News*, *Huffington Post* and *Mashable*) for a period of two months. We ended up with a total of 27,548 tweets and 2,576,853 retweets to these tweets. Figure 1 shows the temporal distribution of the tweets posted by each news media account. Specifically, Figure 1a shows the number of tweets posted per day during the whole period,

whereas Figure 1b shows the number of tweets posted per hour during one week of observation. Focusing in Figure 1a, we observe certain periodicity in the number of tweets posted by our news media per day, seeing that the number of tweets on weekends is lower than on weekdays, and the day with less tweets is Sunday. However, the day when the most tweets are posted varies from one week to the other. In the case of number of tweets per hour (Figure 1b), we also observe a periodic behaviour, where most of the tweets are posted between 12am and 2am. Given that the scale is in *UTC + 0*, the hours in which less tweets are posted correspond to night hours in the US. Finally, focusing on individual news media sources, there is clearly a gap between the number of tweets posted by *Huffington Post*, *Mashable* and *The New York Times* with respect to *BBC* and *CNN Breaking News*. This is because we consider the accounts of *BBC* and *CNN* that only post *breaking news* (news very popular that receive high engagement from their followers), whereas for the others, we consider their regular accounts.


3.1.1 Time-dependence of retweets

So far, we have described the dataset in terms of tweets posted by news media accounts, but now we focus on the retweets to these tweets. And specifically, in the number of retweets that every tweet has received (Figure 2a) and the time between the publication of tweets and their retweets (Figure 2b). Although the majority of tweets posted by our news media receive some retweets, there are huge differences in the number of retweets received by them. For instance, *BBC* and *CNN* are the ones that receive more engagement from their users, mainly because they only post breaking news (popular news). On the contrary, *Huffington Post*, apart from being the news media that publishes the most tweets, is the one whose tweets receive the least engagement (less retweets) from their followers.

Figure 2b contains the temporal distribution of the average number of retweets per tweet in the hour after their publication. Here we see that, independent of the news media considered, (i) tweets get more engagement in the second and third minute after their publication and (ii) since the second minute, the number of retweets achieved fits a power law distribution. Therefore, what is clear is that the most recent tweets, i.e. those in the top-positions of the timeline, are the most exposed to the users and have more possibilities of being retweeted. Finally, Figure 2c shows that the scattering, and therefore the variance, of the number of retweets received between 2 and 9 minutes after publication is higher than after 10 minutes, with the number decreasing significantly after 38 minutes from the tweet publication. This shows that, with the exception of the retweets produced in the first minute after publication, the variance of the number of retweets received per minute decreases over time.

3.1.2 Temporal comparison with other platforms


In our system, we only consider tweets posted by news media and therefore most of them contain news with limited lifetimes. For this reason, the time difference between a retweet and its original tweet is very less, and the majority of retweets occur within the first hour of publication of the original tweet. However, when considering the whole of Twitter, only half of the retweets are made during this first hour, whereas the rest can even happen a month after the publication of the original tweet [15]. Although we



(a) # of tweets posted per day during the 2 months of observation.

(b) # of tweets posted per hour during 2 days of observation.

Figure 1: Temporal distribution of tweets publication.



(a) CCDF of the # of retweets received per tweet for each one of the news accounts considered.

(b) Average # of retweets received per tweet after $time$ minutes of its publication (split by news account)

(c) Scattering of the # of retweets received per tweet after $time$ minutes of its publication for all the tweets in the dataset.

Figure 2: (Temporal) distribution of retweets per tweet.

selected Twitter for our study, other systems providers of content present a similar tendency in terms of the drop of the attention received with the loss of novelty. This is the case of the interactive website *digg.com*, in which the attention received by the stories that users discover from the Internet and upload to the platform decay with their loss of novelty [21]. However, this reduction, at least for popular stories, is not as noticeable as in the case of Twitter, since this platform (*digg*) prioritizes the information to display taking into account their novelty, but also their popularity. So, the more times popular items are displayed, the more prone they are to get users' attention. As a conclusion, in

Twitter, and especially when considering news, the decay is faster than in other platforms, and therefore the interval between discrete times in which our system can update the display should be chosen to be short.

3.1.3 Conditional variance

We also calculated the conditional variance of the number of retweets received after t minutes from the publication of their original tweets. This is the variance in terms of number of retweets received after t minutes from publication by those tweets that had received the same number of retweets after $t - 1$ minutes. Our results revealed variance values that were

larger than zero for all the different values of retweet counts in t .

3.2 Setting up the algorithm parameters

3.2.1 States and reward

In Twitter, we can track a certain set of properties for each tweet, such as age, number of retweets, favourites, etc. We consider that the properties that define the state of each item -tweet- at each instant t are its *novelty* -time since publication- and *popularity* -understanding this as the *number of retweets* that it has received-. In order to have a finite set of states E , we discretize the possible values of novelty and number of retweets, ending up with 10 different values for *novelty* and 10 for *popularity*. Hence each state can be represented as a 2-vector $(n, p) \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$, where n is the novelty -time since publication- and p is the popularity -number of retweets received-. In addition to those 100 states, we also consider the state 0, the “unknown state”. Each item initially starts in this state and also ends on it, with this state serving as both the sink and the source.

In order to set the reward and the values of the properties that define each state we look at the *novelty* and *popularity* of the tweets posted during the first month of observation. Figure 3a shows the average number of retweets per minute that tweets receive during the hour immediately after their publication, whereas Figure 3b represents the CCDF (Complementary Cumulative Distribution Function) of the number of retweets received by those tweets posted during the *first month* of observation. Focusing on Figure 3a, we observe that, on average, tweets receive the majority of retweets in the first minutes after publication and it is also in these first minutes when there are the highest differences between the average number of retweets received in one interval with respect to the others. Therefore, we decided to set the limits between the different intervals that define the state to

$$\lim_n = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 20, 60\}. \quad (10)$$

So, the state of novelty $i \in n$ contains the tweets that were posted between $\lim_n[i]$ and $\lim_n[i + 1] - 1$ minutes before the current time of observation.

Focusing on the popularity of the tweets -*number of retweets received*- (Figure 3b), we observe that the the number of retweets per tweet is distributed according to a power law distribution, where the majority of the tweets receive less than 100 retweets whereas a very small percentage of tweets is retweeted more than 1000 times. In order to set the popularity of the states, we split the tweets, sorted according to the times they are retweeted, into equal sized subsets. So, the limits between the different intervals that define the state are

$$\lim_p = \{0, 1, 19, 25, 32, 39, 48, 61, 82, 131, \infty\}. \quad (11)$$

So, the state of popularity $j \in p$ contains the tweets that have been retweeted between $\lim_p[j]$ and $\lim_p[j + 1] - 1$ times before the current time of observation.

Finally, we reward the tweets in proportion to the number of retweets they receive and their novelty. So, we set the reward of each state to

$$\begin{cases} r(n, p) &= r_n * r_p, \\ r(0) &= 0 \end{cases} \quad (12)$$

where the r_n and r_p are the normalized average number of retweets per interval. That is, the average number of retweets received between $\lim_n[i]$ and $\lim_n[i + 1] - 1$ minutes after publication in the case of novelty, and the average number of total retweets received by those tweets that have received between $\lim_p[i]$ and $\lim_p[i + 1] - 1$ retweets in the case of popularity. This results in

$$\begin{aligned} r_n &= \{0.28, 1, 0.92, 0.79, 0.63, 0.51, 0.43, 0.37, 0.21, 0.07\}, \\ r_p &= \{0.01, 0.07, 0.11, 0.16, 0.19, 0.24, 0.28, 0.34, 0.44, 1\}. \end{aligned} \quad (13)$$

Please, note that the reward when $p = 1$ is not zero but, in order to conserve the reward of the novelty in $r(n, 1)/n \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$, we consider that the average number of retweets in this set is 1.

3.2.2 Transition probabilities

As mentioned, we assume that the state of each item changes according to a Markov process independent of the state of other items, with transition probabilities $\{P_{ij}^1 : i, j \in E\}$ if the item is on the top list, and $\{P_{ij}^0 : i, j \in E\}$ if it is not. In order to empirically calculate these transition probabilities we consider all the tweets posted during the first month of observation. Assuming that all the items are on the top list (all of them are displayed), we define $\{P_{ij}^1 : i, j \in E\}$ as


$$P_{ij}^1 = \frac{|I_j(t + 1)|}{|I_i(t)|} \quad (14)$$

where $I_i(t)$ is the set of items in state i at time t and $I_j(t + 1)$ the set of items in state j at $t + 1$ that transitioned to this state from state i . Finally, we fix $\epsilon_i = 0.1$ for all $i \in E$, which expresses the fact that displaying an item on the top list accelerates its transition speed by ten times.


3.3 Solution

The G index rankings of the 101 states are calculated using the Bertsimas-Niño-Mora heuristic described in Section 3. The results are shown in Figure 4. As can be seen from the figure, state (2, 10) has the largest G index, state (10, 3) the second-largest, and so on. As mentioned, the absolute value of the indices are not as important as their relative orders: items -tweets- should be displayed according to the relative order of the indices of their states.

The result is by no means trivial. For example, the top state (1) is not the most novel but it is the most popular. On the other hand (6) is less popular but more novel than (7). Also, the fact that the algorithm gives high index values to potentially valuable states means that the unknown state which gives no reward should have higher display priority than other states with positive reward. Finally, note also that the influence of the popularity in the output is higher than the novelty, which supports our premise that the current novelty itself doesn't maximise the expected value (informativeness) for the user.



(a) Average # of retweets received per tweet after $time$ minutes of its publication for all the tweets posted during the first month of observation.



(b) CCDF of the # of retweets received per tweet for all the tweets posted during the first month of observation.

Figure 3: (Temporal) distribution of retweets per tweet during the first month of observation.

4. EVALUATION

Now we validate if the states, transition probabilities and, mainly, the ranking of tweets obtained with the *Huberman-Wu* algorithm guarantee that (i) tweets are arranged according to the utility they will have for the users and (ii) the attention they will receive. With this aim, we validate our approach by measuring the degree of similarity between the ranking of tweets according to their expected utility in $t + 1$ (Section 3.3) and that according to their actual utility (reward) in $t + 1$. To perform this evaluation, we consider all the tweets and retweets produced during the last month of observation. In the rest of this section, we detail the experiments conducted and the results obtained.

4.1 Experiments

Assuming that the system (Twitter) could update users' timelines at discrete time intervals of one minute and that the *active* tweets in each instant (minute) t —those that could be displayed—are those that have been posted in the last hour, the steps of our validation (in each instant t) are the following:

1. Find out the state of each tweet in t according to its novelty and popularity and rank the tweets according to their *expected utility* in $t + 1$ (ranking of their states obtained in Section 3.3).
2. Find out the state of each tweet in $t + 1$ and rank them according to their *actual utility* in $t + 1$.
3. Measure the similarity between the rankings in (1) and (2).

4.2 Results

In order to measure the level of similarity between the rankings in (1) and (2), we used the well known *Normalized Discounted Cumulative Gain* ($nDCG$) [13], considering the *actual utility* of the tweets in $t + 1$ as an indicator of their relevance. The formula of $nDCG$ is given in Equation 15:

$$nDCG = \frac{1}{Z} \sum_{p=1}^k \frac{2^{s(p)} - 1}{\log(1 + p)} \quad (15)$$

where k is the number of *active items* in t , $s(p)$ is the function that represents the reward (actual utility) in $t + 1$ given to the tweet at position p , and Z is a normalization factor derived from the perfect ranking of tweets that yields a maximum $nDCG$ value of 1.

Using the utility in $t + 1$ as the ground truth for computing the $nDCG$ for each time t ($nDCG = 1$ means perfect ranking), the $nDCG$ evaluation result for each instant t is shown in Figure 5. Specifically, Figure 5a contains the distribution of the different values of $nDCG$ over time, whereas Figure 5b shows the CCDF of the number of active tweets per minute. The average of $nDCG$ obtained is close to 1 (0.97), which means that our method predicts with high accuracy the expected utility of the tweets. On the other hand, the (Pearson) correlation between the $nDCG$ and the number of *active* tweets in each instant is almost zero (-0.06), which shows the independence of the predictive power of our methodology with the number of items (tweets) to rank.

4.3 Validation considering users' attention

Now, we check if the ranking obtained with our methodology is such that, in each time t , tweets are ranked according to the attention that they will receive in the next time interval $[t, t + 1]$. To this aim, and given that the number






Figure 4: The 101 states ranked by their G indices, ranked from highest to lowest, plus the state 0 (ranked in position 93).



(a) $nDCG$ values obtained during each instant of decision.



(b) CCDF of the # of active tweets during each instant of decision.

Figure 5: Results.

Table 1: $nDCG$ values considering attention

	average	std. dev.
# RT	0.76	0.21
# RT + # replies	0.76	0.21
# RT + # replies + # favourites	0.69	0.25

of users that read every tweet is not available, we approximate the attention of every tweet by the number of retweets, replies and favourites that they have received in the interval $[t, t + 1]$. Although this is not accurate, it gives us an approximation of the number of views received by a tweet, since the more people read it, the more prone it is to receive retweets (replies or favourites).

Table 1 contains the average and standard deviation of the $nDCG$ for all the instants of decision, considering the attention received by the tweets (number of retweets, replies and favourites) in $[t, t + 1]$ as indicators of their relevance. The average $nDCG$ obtained was 0.76 in the case of only considering the number of retweets received (# RT in the figure) as indicator of attention, with a standard deviation of 0.21. Similar values are obtained when the attention is approximated by the sum of retweets and replies received. However, the results are worse in the case of considering also the number of favourites (lower average and higher standard deviation). This is because the temporal distribution of the number of favourites in our dataset is not accurate since we can only retrieve, from the Twitter API, tweets, retweets and replies, but the number of favourites of a tweet must be inferred from the aforementioned entities.

4.4 Comparison between our ranking and those according to novelty and popularity

Finally, we compare the arrangement of tweets obtained

with our methodology with those according to (i) novelty and (ii) popularity (expressed as number of retweets received) of the tweets in time t . To this aim, we compute the $nDCG$ considering the utility and the attention received in the next time interval as indicators of their relevance. Results are displayed in Table 2.

The results show that the ranking of tweets according to the expected utility in $t + 1$ obtained with our methodology presents a higher level of similarity with that according to the real utility in $t + 1$ (ground truth) than the ranking according to the novelty and, specially, the one according to the popularity of the tweets in t . Also, and what it is more important, our ranking of tweets is more similar to the attention received in the interval $[t, t + 1]$ than the rankings according to novelty and to popularity. As a conclusion, top-tweets obtained by our methodology are, in average, tweets with higher utility for the users in $t + 1$ and will receive more attention in the interval $[t, t + 1]$ than the top-tweets in the rankings according to (i) novelty and (ii) popularity in t .

As justification for the slight improvement achieved by the arrangement of tweets proposed by our model with respect to the one based on novelty, consider that, in the case of novelty, users see the tweets in such a way that the most recent are “the most exposed” (in the top positions) to the users and therefore are more likely to be seen (and therefore retweeted) in the next time interval.

5. ANALYSIS CONSIDERING ONLY PEAK HOURS

We repeated the experiment but considering only those tweets and retweets posted during daily hours in the US - peak hours-, hours in which the majority of the tweets in our dataset were posted (see Section 3.1). We fixed the states and the transition probabilities taking into account

Table 2: Comparison between our arrangement of tweets and the arrangement according to (i) novelty and (ii) popularity of tweets.

	Our method		novelty		popularity	
	average	std. dev.	average	std. dev.	average	std. dev.
Utility	0.97	0.05	0.85	0.09	0.68	0.11
# RT	0.76	0.21	0.66	0.21	0.46	0.21
# RT + # replies	0.76	0.21	0.67	0.21	0.45	0.21
# RT + # replies + # favourites	0.69	0.25	0.62	0.24	0.40	0.21

only those tweets posted between 12am and 2am of the first month of observation, and validated our proposal with those tweets posted between 12am and 2am of the second month. We kept the limits between the states according to their novelty ($lim_n = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 20, 60\}$), but the limits according to their popularity were recalculated applying the methodology explained in Section 3.2.1 over the subset of tweets ($lim_p = \{0, 1, 18, 25, 31, 38, 47, 59, 80, 133, \infty\}$). This also resulted in new rewards for the states: $r_n = \{0.28, 1, 0.90, 0.76, 0.61, 0.49, 0.41, 0.35, 0.19, 0.06\}$ and $r_p = \{0.01, 0.05, 0.09, 0.13, 0.17, 0.20, 0.24, 0.31, 0.39, 1\}$, in which the difference between the reward of the *top* state and the other states has been increased with respect to that in the initial analysis (Section 3.2.1).

Results of this analysis displayed in Table 3 show that the average value of the *nDCG* for the expected utility when considering only peak hours is slightly lower (0.96) than when considering all the tweets posted during the whole day (0.97). The reason could be that the more tweets used to define the states and calculate the transition probabilities, the more accurate the parameters are and therefore better accuracy in the prediction. On the contrary, the average *nDCG* for the attention received when considering only peak hours is slightly better than when considering the tweets posted during the whole day (0.77 in the case of peak hours in comparison with the 0.76 when considering all the tweets). Similar results were obtained when measuring the *nDCG* considering the rankings according to novelty and according to popularity during peak hours, with the ranking according to popularity being the one that achieves the highest improvement when we consider all the tweets. Finally, it is noticeable that there is an increase in the average value of *nDCG* when tweets posted during peak hours are ranked according to their popularity as compared to its value when all the tweets are ranked.

6. RELATED WORK

Since the emergence of the Web, the exponential growth of information has lead to the development of algorithms to rank the information that users receive when make a particular query. Ranking [2], therefore, is at the core of the information retrieval (IR) scientific discipline behind search engines as *Google* and *Yahoo! Search*. Although ranking has been studied for decades, a recent trend deals with applying machine learning techniques to learn to rank functions automatically: *Learning to rank* [16]. However, and independently of the algorithm used to rank in IR, the problem we aim to solve is slightly different to the typical problem in IR: on Twitter, like on other social media platforms, tweets are displayed in users’ timelines without them have to specify any explicit query in order to obtain relevant tweets from the system.

Closer to our scenario it is the problem of ranking social streams. Users’ timelines on social media platforms as *Facebook*, *Twitter* or *LinkedIn* have recently been the natural field of application of algorithms that rank or prioritize content according to the relevance for their owners [14, 8, 5, 11]. As an outstanding example, the well known Facebook *EdgeRank* [14] ranks the items in user’s *news feed* according to a function that depends on the novelty of the item, the type of item and the strength of the relationship between the user and the creator of the item. Independently of the ranking algorithm, all of these scenarios have in common that (i) items are sorted according to their relevance to the user instead of their relevance to an explicit query as in the traditional IR and (ii) the ranking of the items depends on the owner of the timeline (user-dependent ranking). In order to select or *recommend* the most attractive and relevant content to the user, these solutions often use personalization techniques based mainly on matching users’ profiles (previously obtained) with items’ content [7, 19, 9, 1, 6]. Contrarily, our proposal aims to provide a ranking of tweets that maximise the utility or relevance for *all* Twitter users, and not only for a specific user or group of users.

Moreover, although in our model we consider retweets as both indicators of popularity and attention, the aim of our study is not to predict if a given user will retweet a tweet or how many times a tweet will be retweeted [4, 22, 17], but to automatically decide, on the basis of its novelty and popularity, if a tweet should be shown in a top or a bottom position of users’ timelines.

7. CONCLUSION

Social media platforms, and in particular the microblogging service Twitter, are examples of competing attention environments whose subscribers have to decide what content to prioritize to their followers in order to get the most attention. The fact that Twitter displays tweets in decreasing order of publication limits the capability of tweet promotion, since the more recent tweets are the ones to be displayed in the top positions and the time they remain in these positions only depends on the number of tweets posted in the following minutes. Therefore, tweets that users’ find valuable are frequently hidden. As a solution, we customized the *Huberman-Wu* algorithm [12] to select the optimal arrangement of tweets that maximises utility for the users. We experimentally validated this method using real data from Twitter. Our results confirmed the suitability of using the aforementioned algorithm in order to, not only arrange the tweets according to their utility for the users, but also maximize users’ attention.

So far, we have focused our study on Twitter, but the extension of this mechanism to other social news aggregators, such as *digg.com* or *reddit.com* is straightforward. Also,

Table 3: Comparison between our arrangement of tweets and the arrangement according to (i) novelty and (ii) popularity of tweets (peak hours).

	Our method		novelty		popularity	
	average	std. dev.	average	std. dev.	average	std. dev.
Utility	0.96	0.06	0.82	0.10	0.83	0.11
# RT	0.77	0.22	0.68	0.22	0.58	0.23
# RT + # replies	0.77	0.23	0.67	0.22	0.58	0.23
# RT + # replies + # favourites	0.72	0.25	0.64	0.24	0.54	0.24

although we only consider a particular kind of tweets (those that contain news), this methodology could also be extended to other kinds of tweets such as promotional campaigns and status updates from users. This would require dealing with different temporal patterns of retweet behaviour. And, the more patterns exist, the more different systems must be considered (in terms of definition of states and calculation of their transition probabilities).

Finally, in our model we have taken into account the retweets made by all Twitter users independent of the specific interests of the users we display the items to. As future work, we plan to add personalization to our methodology to select the tweets to display. That is, different arrangement of tweets will be displayed to different users –groups of users– on the basis of their interests, the interests of their friends or the interests of users with similar interests (users that historically have retweeted –replied, marked as favourites– the same tweets) as the given user.

8. ACKNOWLEDGMENTS

Work partially funded by the Spanish Ministry of Economy and Competitiveness (EEBB-I-14-07966, TEC2013-47665-C4-3-R); the European Regional Development Fund and the Galician Regional Government under agreement for funding the AtlantTIC Research Center; and the Spanish Government and the European Regional Development Fund under project TACTICA.

9. REFERENCES

- [1] F. Abel, Q. Gao, G.-J. Houben, and K. Tao. Analyzing user modeling on twitter for personalized news recommendations. In *User Modeling, Adaption and Personalization*, pages 1–12. Springer, 2011.
- [2] R. Baeza-Yates, B. Ribeiro-Neto, et al. *Modern information retrieval*, volume 463. ACM press New York, 1999.
- [3] D. Bertsimas and J. Niño-Mora. Conservation laws, extended polymatroids and multiarmed bandit problems; a polyhedral approach to indexable systems. *Mathematics of Operations Research*, 21(2):257–306, 1996.
- [4] D. Boyd, S. Golder, and G. Lotan. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In *System Sciences (HICSS), 2010 43rd Hawaii International Conference on*, pages 1–10. IEEE, 2010.
- [5] J. Chen, R. Nairn, L. Nelson, M. Bernstein, and E. Chi. Short and tweet: experiments on recommending content from information streams. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1185–1194. ACM, 2010.
- [6] G. De Francisci Morales, A. Gionis, and C. Lucchese. From chatter to headlines: harnessing the real-time web for personalized news recommendation. In *Proceedings of the fifth ACM international conference on Web search and data mining*, pages 153–162. ACM, 2012.
- [7] E. Diaz-Aviles, L. Drumond, L. Schmidt-Thieme, and W. Nejdl. Real-time top-n recommendation in social streams. In *Proceedings of the sixth ACM conference on Recommender systems*, pages 59–66. ACM, 2012.
- [8] Y. Duan, L. Jiang, T. Qin, M. Zhou, and H.-Y. Shum. An empirical study on learning to rank of tweets. In *Proceedings of the 23rd International Conference on Computational Linguistics*, pages 295–303. Association for Computational Linguistics, 2010.
- [9] W. Feng and J. Wang. Retweet or not?: personalized tweet re-ranking. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 577–586. ACM, 2013.
- [10] J. C. Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 148–177, 1979.
- [11] L. Hong, R. Bekkerman, J. Adler, and B. D. Davison. Learning to rank social update streams. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 651–660. ACM, 2012.
- [12] B. A. Huberman and F. Wu. The economics of attention: maximizing user value in information-rich environments. *Advances in Complex Systems*, 11(04):487–496, 2008.
- [13] K. Järvelin and J. Kekäläinen. Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems (TOIS)*, 20(4):422–446, 2002.
- [14] J. Kincaid. Edgerank: The secret sauce that makes facebook’s news feed tick. *TechCrunch*, April, 2010.
- [15] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In *Proceedings of the 19th international conference on World wide web*, pages 591–600. ACM, 2010.
- [16] T.-Y. Liu. Learning to rank for information retrieval. *Foundations and Trends in Information Retrieval*, 3(3):225–331, 2009.
- [17] H.-K. Peng, J. Zhu, D. Piao, R. Yan, and Y. Zhang. Retweet modeling using conditional random fields. In *Data Mining Workshops (ICDMW), 2011 IEEE 11th International Conference on*, pages 336–343. IEEE, 2011.
- [18] P. Resnick and H. R. Varian. Recommender systems.

Communications of the ACM, 40(3):56–58, 1997.

- [19] I. Uysal and W. B. Croft. User oriented tweet ranking: a filtering approach to microblogs. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, pages 2261–2264. ACM, 2011.
- [20] P. Whittle. Restless bandits: Activity allocation in a changing world. *Journal of applied probability*, pages 287–298, 1988.
- [21] F. Wu and B. A. Huberman. Novelty and collective attention. *Proceedings of the National Academy of Sciences*, 104(45):17599–17601, 2007.
- [22] Z. Yang, J. Guo, K. Cai, J. Tang, J. Li, L. Zhang, and Z. Su. Understanding retweeting behaviors in social networks. In *Proceedings of the 19th ACM international conference on Information and knowledge management*, pages 1633–1636. ACM, 2010.